IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re The Application of: John B. Duffie III et al.	)	
Serial No.: 10/087,376	)	Examiner: Daftuar, Saket K.
Filed: March 1, 2002	)	Art Unit: 2151
For: METHOD TO OPTIMIZE THE LOAD BALANCING OF PARALLEL CO-PROCESSORS	)	

Cesari and McKenna, LLP
88 Black Falcon Avenue
Boston, MA 02210
February 5, 2008

By EFS

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

APPEAL BRIEF

In response to the Notice of Appeal filed Dec. 5, 2007, the Applicants hereby submits this Appeal Brief.

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST

The real party in interest is Cisco Technology, Inc. by an Assignment recorded at reel 012672 frame 0284.

II. RELATED APPEALS AND INTERFERENCES

The Applicants and the Applicants' legal representatives know of no related appeals or interferences that will directly affect, be directly affected by, or have a bearing on the Board's decision in the present appeal.

III. STATUS OF CLAIMS

The status of the claims is:

A. Total Number of Claims in Application

Claims 1-35 stand pending in the Application.

B. Status of All the Claims

No claims stand cancelled.

Claims 1-35 stand rejected.

No claims are withdrawn from consideration on appeal.

No claims stand allowed.

C. Claims on Appeal

Claims 1-35 are currently being appealed.

IV. STATUS OF AMENDMENTS

No Amendments were filed subsequent to the Final Office Action dated Sept. 5, 2007.

V. SUMMARY OF CLAIMED SUBJECT MATTER

This summary describes exemplary embodiments of the invention.

Independent claim 1 is directed to a method for selecting a coprocessor from a plurality of coprocessors to process a packet. See Application page 9, lines 1-3. A size of a packet is determined. See page 4, lines 17-18 and page 10, lines 9-11. In response to the size of the packet, a cost associated with the with the packet is determined. See page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5. The cost represents a load associated with processing the packet. See page 10, lines 8-9. An anticipated load for each coprocessor in the plurality of coprocessors is determined using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. A coprocessor from the plurality of coprocessors is selected based upon the anticipated load. See page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 17 is directed to an apparatus for selecting a coprocessor from a plurality of coprocessors to process a packet. See page 9, lines 1-3. A memory (see page 7, lines 24-30 and Fig. 3, 340) contains one or more software routines configured to determine a size of a packet (see page 4, lines 17-18, page 10, lines 9-11) and to determine a cost associated with the packet in response to the size of the packet (see page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5). The cost represents a load associated with processing the packet. See page 10, lines 8-9. A processor (see page 7, lines 4-11 and Fig. 3, 320) is configured to execute the software routines to determine an anticipated load for each coprocessor in the plurality of coprocessors using the cost (see page 10, line 26 to page 11, lines 6 and Fig. 6, 604) and to select a coprocessor from the plurality of coprocessors based on the anticipated load (see page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618).

Independent claim 21 is directed to an intermediate device configured to select a coprocessor from a plurality of coprocessors to process a packet. See page 9, lines 1-3. The device includes means for determining a size of the packet (see page 4, lines 17-18, page 10, lines 9-11) and for determining a cost associated with the packet in response to the size of the packet (see page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5).

The cost represents a load associated with processing the packet. See page 10, lines 8-9. The device includes means for determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. Further, the device includes means for selecting a coprocessor based upon the anticipated load. See page 11, lines 7-11 and page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 22 is directed to a computer readable media comprising computer executable instruction for execution in a processor for selecting a coprocessor from a plurality of coprocessors to process a packet. See page 9, lines 1-3. The instruction are provided for determining a size of the packet (see page 4, lines 17-18, page 10, lines 9-11) and determining a cost associated with the packet in response to the size of the packet (see page 9, lines 15 to page 10, line 13 and Fig. 4, 402 and Fig. 5). The cost represents a load associated with processing the packet. See page 10, lines 8-9. Instructions are also provided for determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. Further, instructions are provided for selecting a coprocessor from the plurality of coprocessors based upon the anticipated load. See page 11, lines 7-11 and page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 23 is directed to a method for selecting a processor for processing a packet. See page 9, lines 1-3. A size of a packet is determined. See page 4, lines 17-18 and page 10, lines 9-11. A cost associated with the packet of that size is determined. See page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5. The cost represents a load associated with processing the packet. See page 10, lines 8-9. An anticipated load for the processor is determined using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. The processor is selected based upon the anticipated load. See page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 26 is directed to a method for selecting a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet. See page 9, lines 1-3. A cumulative load for each processor is determined, the cumulative

load representing load due to packets currently awaiting processing. See page 10, line 28 to page 11, line 4. A cost for processing the received packet at each coprocessor is determined, at least in part, in response to the size of the received packet and a processing rate of that coprocessor. See page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5. The cumulative load is combined with the cost at each coprocessor to create an anticipated load for each coprocessor. See page 10, line 26 to page 11, line 6. The anticipated loads of all the coprocessors are compared. See page 11, lines 6-11. In response to the comparison, a particular coprocessor of the plurality of coprocessors is selected to perform the processing operation on the received packet. See page 12, line 15-21 and Fig. 6, 616, and 618.

Independent claim 31 is directed to an apparatus to select a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet. See page 9, lines 1-3. A plurality of queues are configured to store packets currently awaiting processing, each queue associated with one of the coprocessors. See page 7, lines 19-23 and Fig. 3, 323a, 323b, 323c. Each queue is associated with a cumulative load that represents a load to process packets in that queue. See page 10, line 28 to page 11, line 4. A data structure is configured to store processing rates, each processing rate associated with one of the coprocessors. See col. 9, lines 23-30 and Fig. 3, 345. A processor is configured to determine a size of the received packet (see page 4, lines 17-18 and page 10, lines 9-11), and in response to the size of the received packet, and the processing rate of each coprocessor, determine a cost to perform a processing operation on the received packet of each coprocessor (see page 9, line 15 to page 10, line 13 and Fig. 4, 402 and Fig. 5). The processor is further configured to combine the cost at each coprocessor with the cumulative load at that coprocessor to create an anticipated load at each coprocessor (see page 10, line 26 to page 11, line 6), and to select a particular coprocessor to perform the processing operation on the received packet in response to comparison of the anticipated load at each processor (see page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Whether claims 1-9, 13-14, 17, 21-27, 31 and 32 are unpatentable under 35 U.S.C. §102(e) over U. S. Patent No. 6,370,560 to Robertazzi et al. (hereinafter "Robertazzi"), where the reference fails to teach several elements of the claims.

B. Whether claims 10-12, 15-16, 18-20, 28-29, 30, 33, 34 and 35 are unpatentable under 35 U.S.C. §103(a) over Robertazzi in view of U. S. Patent No. 6,587,866 to Modi et al. (hereinafter "Modi"), where the references fail to teach several elements of the claims.

C. Whether claims 30 and 35 are unpatentable under 35 U.S.C. §103(a) over Robertazzi and Modi in view of U. S. Patent No. 6,065,046 to Fienberg et al. (hereinafter "Fienberg"), where the references fail to teach several elements of the claims.

VII. ARGUMENT

A. Grouping of Claims

For each ground of rejection that the Applicant contests herein that applies to more than one claim, such additional claims, to the extent separately identified and argued below, do not stand or fall together.

B. Legal Standard

In rejecting claims under 35 U.S.C. §102, the Examiner bears the initial burden of presenting a prima facie case of anticipation. To establish a prima facie case of anticipation, the Examiner must show a single prior art reference discloses each and every element of a claimed invention. See, e.g., RCA Corp. v. Applied Digital Data Systems, Inc., 730 F.2d 1440, 1444 (Fed. Cir. 1984). There can be no differences between the claimed invention and the cited references, as viewed by a person of ordinary skill in the art. Scripps Clinic & Research Foundation v. Genetech Inc., 927 F.2d 1565, 1576 (Fed. Cir. 1991). If a single claim limitation is missing from the prior art reference, then the rejection is improper and should be reversed.

In rejecting claims under 35 U.S.C. §103, the Examiner bears the initial burden of presenting a prima facie case of obviousness. See, e.g., In re Rijckaert, 9 F.3d 1531, 1532 (Fed. Cir. 1993). To establish a prima facie case of obviousness, the Examiner must show all the elements of an invention are found in a combination of prior art references. U.S. Surgical Corp. v. Ethicon, Inc., 103 F.3d 1554, 1564 (Fed. Cir. 1997); Abbott Laboratories v. Sandoz, Inc., 500 F.Supp.2d 846, 851-852 (N.D. Ill. 2007). In coming to this determination, the Examiner should (a) determine the scope and content of the prior art, (b) ascertain the differences between the prior art and the claims at issue, (c) resolving the level of ordinary skill in the pertinent art, and (d) evaluate evidence of secondary considerations. KSR Intern. Co. v. Teleflex Inc., 127 S. Ct. 1727, 1729 (2007) (citing Graham v. John Deere Co. of Kansas City, 383 U.S. 1, 17-18 (1966)). If the rejection can

only be made by resort to speculation, unfounded assumption, or hindsight reconstruction, then the rejection is improper and should be reversed. In re Warner, 379 F.2d 1011, 1017 (CCPA 1967).

C. Rejection of Claims under 35 U.S.C. §102(e) over Robertazzi

1. Claims 1-6, 8-9, 13-14, 17, 21-25

Independent claim 1, representative of claims 1-6, 8-9, 13-14, 17, 21-25, recites (with emphasis added):

> 1. A method for selecting a coprocessor from a plurality of coprocessors to process a packet, the method comprising the steps of:
> determining a size of the packet;
> ***determining a cost associated with the packet*** in response to the size of the packet, ***the cost representing a load associated with processing the packet***;
> ***determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost***; and
> ***selecting the coprocessor from the plurality of coprocessors based on the anticipated load.***

Robertazzi describes a technique for assigning segments of a "divisible load or task" to different processing platforms (for example "a 486 computer, a workstation, and a supercomputer") to minimize a monetary cost measured in dollars and cents. See Robertazzi col. 4, lines 8-21. Specifically, "[e]ach different type of processing platform has an associated resource utilization cost." Id. col. 4, lines 14-15. This resource utilization cost takes into account the "cost per second" of using a processing platform, and the "loads/tasks per second" the computing platform can process. Id. col. 4, lines 18-22 and col. 5, lines 40-42. "For example, if a 486 computing platform has a computing cost of $5/second and a computer speed of two loads/tasks per second, the monetary cost of processing a unit load/task is $2.250." Id. col. 4, lines 28-31. Similarly, "[i]f a minicomputer platform has a computing cost of $10/second and a computing speed of eight loads/tasks per second, the cost of processing a unit load/task is $1.25." Id. col. 4, lines

31-33. Robertazzi tries to assign tasks to the cheapest available processing platforms (Fig. 2A, 205) that can complete the tasks before a finish time. Id. col. 8, lines 34-36.

In separate discussion, Robertazzi states that availability of processing platforms is maintained in a field (Fig. 1C, 183) in a file (Fig. 1C, 150) that indicates whether each processing platform is "busy" or "available." Id. col. 6, lines 41-46.

First, Robertazzi does not teach the Applicant's claimed "*determining a cost associated with the packet ... the cost representing a load associated with processing the packet.*"

While the Applicant claims a "*cost representing a load associated with processing the packet,*" Robertazzi describes a quite different type of "cost," a monetary cost. A monetary cost does not represent a load or burden on a processor, but simply represents some amount of dollars and cents that must be paid in exchange for the right to use a processor.

The difference may be illustrated by way of example. For instance, at page 10, lines 8-13 of the Specification, the Applicant discusses calculating a cost that represents a load, and determines in the example that the cost is 25 milliseconds of processing time. Such description stands in contrast to Robertazzi's discussion at col. 4, lines 28-31, which yields a monetary cost of $2.50 payable in exchange for using a 486 processing platform. As the Applicant's "*cost representing a load associated with processing the packet*" differs considerably from Robertazzi's monetary "cost," the claims rapidly diverge from what Robertazzi discusses.

Second, Robertazzi does not teach the Applicant's claimed "*determining an anticipated load for each coprocessor*" and "*selecting the coprocessor from the plurality of coprocessors based on the anticipated load.*"

The Applicant determines an *anticipated load.* The Applicant defines *anticipated load* at page 9, lines 5-6 of the specification stating "anticipated load is the load a coproc-

essor would incur if it were to process the packet given its current load." The Applicant then selects a coprocessor "*based on the anticipated load*."

Robertazzi lacks any metric akin to the claimed anticipated load, and thus does not select processing platforms based on it. Instead, Robertazzi looks to monetary costs, assigning tasks to processing platforms trying to achieve the cheapest price. Each processing platform in Robertazzi is simply characterized as "available" or "busy." See Robertazzi Fig. 1C, 183. Robertazzi fails to determine how much more busy a processing platform would be if it were to process a task given its current load, nor does Robertazzi base any selection upon this.

Accordingly, for the above reasons, rejection of claims 1-6, 8-9, 13-14, 17, 21-25 should be reversed.

2. Claim 7

Dependent claim 7 recites (with emphasis added):

> 7. The method of claim 1 wherein the step of ***determining an anticipated load further comprises*** the step of:
> ***adding the cost to a cumulative load associated with each coprocessor in the plurality of coprocessors.***

Robertazzi does not teach determining an anticipated load by "*adding the cost to a cumulative load associated with each coprocessor in the plurality of coprocessors.*"

Robertazzi lacks any precise measure of *cumulative load* and thus can not possibly teach calculating anticipated load based on it. Each processing platform in Robertazzi is simply classified as "available" or "busy;" there is no mention of a more specific measure of how available, or how busy, a processor is. See Robertazzi col. 6, lines 44-46 and col. 7, lines 54-55. Without having a measure of cumulative load, Robertazzi would be unable to calculate *anticipated load* by adding cost to "*cumulative load associated with each coprocessor*."

Accordingly, for the above reasons, rejection of claim 7 should be reversed.

2. Claims 26-27 and 31-32

Independent claim 26, representative of claims 26-27 and 31-32, recites (with emphasis added):

> 26. A method for selecting a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet, the method comprising steps of:
>
> ***determining a cumulative load for each coprocessor, the cumulative load representing load due to packets currently awaiting processing at that coprocessor;***
>
> determining a size of the received packet;
>
> determining a cost for processing the received packet at each coprocessor, the cost determined, at least in part, in response to the size of the received packet and a processing rate of that coprocessor;
>
> ***combining the cumulative load and the cost at each coprocessor, to create an anticipated load for each coprocessor***;
>
> ***comparing the anticipated loads of all the coprocessors***; and
>
> selecting, in response to the comparing, a particular coprocessor of the plurality of coprocessors to perform the processing operation on the received packet.

To the extent that the above discussion relating to claims 1 and 7 is relevant to various limitations of claim 26, in the interest of brevity, the reader is referred thereto. The Applicant would like to emphasize, however, that Robertazzi lacks any disclosure of the claimed "*determining a cumulative load for each coprocessor, the cumulative load representing load due to packets currently awaiting processing at that coprocessor*" and "*combining the cumulative load and the cost at each coprocessor, to create an anticipated load for each coprocessor*" and "*comparing the anticipated loads of all the coprocessors.*"

In particular, it should be noted that rather than compare *anticipated loads* to select a coprocessor, Robertazzi compares monetary values. Such reliance on monetary values rather than on load is made clear in Robertazzi. Specifically, Robertazzi states at col. 8, lines 34-41 "[s]tep 205 allocates the cheapest available participating processing platform (based on the cost/unit load tasks)... The load or task will be divided so that the segment will fully utilize the cheapest processor for the entire selected finish time con-

straint." Accordingly, as Robertazzi selects based upon a differing type of comparison, Robertazzi does not teach what the Applicant claims.

Accordingly, for the above reasons, rejection of claims 26-27 and 31-32 should be reversed.

D. Rejection of Claims under 35 U.S.C. §103(a) over Robertazzi in View of Modi

1. Claims 15-16

Dependent claim 15, representative of claims 15 and 16, recites (with emphasis added)

> 15. The method of claim 1 further comprising the step of:
> ***decrementing a cumulative load associated with the selected co-processor.***

Modi discloses a load-balancing scheme for distributing packets among a plurality of server nodes in a clustered processing system. See Modi col. 3, lines 18-22 and col. 7, lines 61-64. The scheme uses one of a plurality of load balancing policy types, included "non-affinity" policy types where packets are distributed to any server node, and "affinity" policy types where packets from a single client are sent to the same server node. Id. col. 7, lines 61 to col. 8, lines 2 and col. 2, lines 42-47.

Neither Robertazzi nor Modi teach the Applicant's claimed *"decrementing a cumulative load associated with the selected coprocessor."*

The Final Office Action of Sept. 5, 2007 simply points to Modi col. 12, lines 54-58 in connection with this claim limitation. The cited portion of Modi reads:

> When an existing TCP connection is terminated, the old node sends a message indicating that the connection has been terminated, and the connection is deleted from the forwarding list (step 1018). When the forwarding list becomes empty the entry for the forwarding list may be removed from the bucket (step 1020).

Modi at col. 12, lines 54-58 has little to do with "*decrementing a cumulative load associated with the selected coprocessor.*" Instead, this portion of Modi deals with terminating a TCP connection and disposing of an empty forwarding list.

Further, review of Modi and Robertazzi finds no other portions that teach what is claimed.

Accordingly, for the above reasons, rejection of claims 15-16 should be reversed.

2. Claims 28 and 33

Dependent claim 28, representative of claims 28 and 33, recites (with emphasis added)

28. The method of claim 26, further comprising the step of:

> ***determining if congestion is present at an output port associated with the received packet, and if congestion is present, selecting a coprocessor with non-minimum anticipated load to perform the processing operation on the received packet.***

Neither Robertazzi nor Modi teach the Applicant's claimed "*determining if congestion is present at an output port associated with the received packet, and if congestion is present, selecting a coprocessor with non-minimum anticipated load to perform the processing operation on the received packet.*"

The Final Office Action of Sept. 5, 2007, at page 8, agrees "Robertazzi is silent about determining if a port associated with the packet is congested" and turns to Modi col. 10, lines 56-58 in relation to the claimed congestion determination. However, the cited portion of Modi reads:

> If the selected bucket has a forwarding list, the system determines if the client IP address and port number are in a forwarding list (step 610).

Modi at col. 10, lines 56-58 has little to do with "*determining if congestion is present at an output port associated with the received packet.*" Instead, Modi at col. 10,

lines 56-58 deals with determining if certain information is present in a forwarding list. Accordingly, such aspect of the claims is not taught by Modi.

Further, neither reference teaches the claimed "*if congestion is present, selecting a coprocessor with non-minimum anticipated load to perform the processing operation on the received packet.*" The Final Office Action of Sept. 5, 2007, at page 8, agrees "Robertazzi is silent about ... selecting coprocessor based on anticipated load not a minimum load," but later, on page 9, disagrees stating "Robertazzi discloses selecting the coprocessor from a group of one or more coprocessors whose anticipated load is not a minimum." Citation is made to col. 4, lines 7-59 and col. 11, line 5 to col. 12, line 39. Such portions of Robertazzi simply describe Robertazzi's selection scheme in response to monetary costs and, as discussed above, lack any mention of *anticipated loads*. Accordingly, such portions may not properly be interpreted as teaching selection of a coprocessor with *non-minimum anticipated load*, if congestion is present on a port.

Accordingly, for the above reasons, rejection of claims 28 and 33 should be reversed.

3. Claim 29 and 34

Dependent claim 29, representative of claims 29 and 34, recites (with emphasis added)

> 29. The method of claim 26, wherein the step of determining a cumulative load for each coprocessor further comprises the step of:
> ***determining, for each coprocessor, sizes of the packets currently awaiting processing at that coprocessor and using the sizes in conjunction with the processing rate of that coprocessor to determine the cumulative load.***

Neither Robertazzi nor Modi teach the Applicant's claimed "*determining, for each coprocessor, sizes of the packets currently awaiting processing at that coprocessor and using the sizes in conjunction with the processing rate of that coprocessor to determine the cumulative load.*"

The Final Office Action of Sept. 5, 2007, at page 11, states "[a]s per claims 28, 29, 33, and 34 they do not teach or further define over the limitations as recited in claims 10-12, and 13-16, respectively. Therefore, claims 27, 28 and 29 are rejected under same scope as recited in claims 10-12 and 13-16, supra." Claim 29 recites quite different limitations than claims 10-12 and 13-16 and thus rejection simply by reference to these claims is improper. For example, claim 29 recites "*determining sizes of the packets currently awaiting processing at that coprocessor*" and using such sizes in part "*to determine the cumulative load*," while claims 10-12 and 13-16 make no mention of sizes of the packets or their use.

Accordingly, for the above reasons, rejection of claims 29 and 34 should be reversed.

4. Claims 10-12, 18-20, 30 and 35

Claims 10-12, 18-20, 30 and 35 are dependent claims that depend from independent claims believed to be allowable for the reasons discussed above. Accordingly, these claims are believed to be allowable at least due to such dependency.

E. Rejection of Claims under 35 U.S.C. §103(a) over Robertazzi and Modi in View of Feinberg

1. Claims 30 and 35

Claims 30 and 35 are dependent claims that depend from independent claims believed to be allowable for the reasons discussed above. Accordingly, these claims are believed to be allowable at least due to such dependency.

VIII. CONCLUSION

The Applicants respectfully submits that the claims are allowable over the art of record. Accordingly, the Applicants requests that the rejection of all claims be reversed.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

James A. Blanchette
Reg. No. 51,477
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500

VIV. CLAIMS APPENDIX

1. (PREVIOUSLY PRESENTED) A method for selecting a coprocessor from a plurality of coprocessors to process a packet, the method comprising steps of:

determining a size of the packet;

determining a cost associated with the packet in response to the size of the packet, the cost representing a load associated with processing the packet;

determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost; and

selecting the coprocessor from the plurality of coprocessors based on the anticipated load.

2. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of determining a cost further comprises the step of:

calculating the cost using a rate associated with processing the packet.

3. (ORIGINAL) The method of claim 2 wherein the rate is stored in a lookup table.

4. (ORIGINAL) The method of claim 2 wherein the step of calculating the cost further comprising the step of:

dividing the packet's size by the rate.

5. (PREVIOUSLY PRESENTED) The method of claim 2 wherein the step of calculating the cost further comprises the step of:

multiplying the packet's size by a multiplicative inverse of the rate.

6. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of determining a cost further comprises the step of:

applying the packet's size to a lookup table containing one or more cost values to determine the cost.

7. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of determining an anticipated load further comprises the step of:

adding the cost to a cumulative load associated with each coprocessor in the plurality of coprocessors.

8. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of selecting the coprocessor further comprises the step of:

selecting the coprocessor from a group of one or more coprocessors whose anticipated load is a minimum load.

9. (ORIGINAL) The method of claim 8 wherein the coprocessor is selected using a scheduling algorithm.

10. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of selecting the coprocessor further comprises the step of:

determining if a port associated with the packet is congested.

11. (PREVIOUSLY PRESENTED) The method of claim 10 wherein the step of selecting the coprocessor further comprises the step of:

selecting the coprocessor from a group of one or more coprocessors whose anticipated load is not a minimum load.

12. (PREVIOUSLY PRESENTED) The method of claim 10 wherein the step of selecting the coprocessor further comprises the step of:

selecting the coprocessor from a group of one or more coprocessors whose anticipated load is a minimum load.

13. (ORIGINAL) The method of claim 1 further comprising the step of:

incrementing a cumulative load associated with the selected coprocessor.

14. (PREVIOUSLY PRESENTED) The method of claim 13 wherein the step of incrementing a cumulative load further comprises the step of:

adding the cost to the cumulative load.

15. (ORIGINAL) The method of claim 1 further comprising the step of:

decrementing a cumulative load associated with the selected coprocessor.

16. (PREVIOUSLY PRESENTED) The method of claim 15 wherein the step of decrementing a cumulative load further comprises the step of:

subtracting the cost from the cumulative load.

17. (PREVIOUSLY PRESENTED) An apparatus for selecting a coprocessor from a plurality of coprocessors to process a packet, the apparatus comprising:

a memory containing one or more software routines, including a software routine configured to determine a size of the packet, and to determine a cost associated with the packet in response to the size of the packet, the cost representing a load associated with processing the packet; and

a processor configured to execute the software routines to determine an anticipated load for each coprocessor in the plurality of coprocessors using the cost and to select the coprocessor from the plurality of coprocessors based on the anticipated load.

18. (ORIGINAL) The apparatus of claim 17 further comprising:

a data structure;

wherein the cost is determined using information contained in the data structure.

19. (ORIGINAL) The apparatus of claim 18 wherein the information contained in the data structure includes the cost.

20. (ORIGINAL) The apparatus of claim 18 wherein the information contained in the data structure includes a rate the coprocessor can process the packet.

21. (PREVIOUSLY PRESENTED) An intermediate device configured to select a coprocessor from a plurality of coprocessors to process a packet, the intermediate device comprising:

means for determining a size of the packet, and for determining a cost associated with the packet in response to the size of the packet, the cost representing a load associated with processing the packet;

means for determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost; and

means for selecting the coprocessor based on the anticipated load.

22. (PREVIOUSLY PRESENTED) A computer readable media comprising computer executable instructions for execution in a processor for selecting a coprocessor from a plurality of coprocessors to process a packet, the instructions for:

determining a size of the packet, and determining a cost associated with the packet in response to the size of the packet, the cost representing a load associated with processing the packet;

determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost; and

selecting the coprocessor from the plurality of coprocessors based on the anticipated load.

23. (PREVIOUSLY PRESENTED) A method for selecting a processor for processing a packet, the method comprising steps of:

determining a size of the packet;

determining a cost associated with the packet of that size, the cost representing a load associated with processing the packet;

determining an anticipated load for the processor using the cost of the packet if processed by the processor;; and

selecting the processor based on the anticipated load.

24. (PREVIOUSLY PRESENTED) The method of claim 23 wherein the step of determining a cost comprises the step of:

calculating the cost using a rate associated with processing of the packet; and

wherein the rate is stored in a lookup table.

25. (PREVIOUSLY PRESENTED) The method of claim 23 wherein the step of determining a cost further comprises the step of:

applying the size of the packet to a lookup table containing cost values associated with particular sizes.

26. (PREVIOUSLY PRESENTED) A method for selecting a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet, the method comprising steps of:

determining a cumulative load for each coprocessor, the cumulative load representing load due to packets currently awaiting processing at that coprocessor;

determining a size of the received packet;

determining a cost for processing the received packet at each coprocessor, the cost determined, at least in part, in response to the size of the received packet and a processing rate of that coprocessor;

combining the cumulative load and the cost at each coprocessor, to create an anticipated load for each coprocessor;

comparing the anticipated loads of all the coprocessors; and

selecting, in response to the comparing, a particular coprocessor of the plurality of coprocessors to perform the processing operation on the received packet.

27. (PREVIOUSLY PRESENTED) The method of claim 26, wherein the step of selecting further comprises the step of:

selecting a coprocessor with minimum anticipated load to perform the processing operation on the received packet.

28. (PREVIOUSLY PRESENTED) The method of claim 26, further comprising the step of:

determining if congestion is present at an output port associated with the received packet, and if congestion is present, selecting a coprocessor with non-minimum anticipated load to perform the processing operation on the received packet.

29. (PREVIOUSLY PRESENTED) The method of claim 26, wherein the step of determining a cumulative load for each coprocessor further comprises the step of:

determining, for each coprocessor, sizes of the packets currently awaiting processing at that coprocessor and using the sizes in conjunction with the processing rate of that coprocessor to determine the cumulative load.

30. (PREVIOUSLY PRESENTED) The method of claim 26 wherein the processing operation is an encryption operation.

31. (PREVIOUSLY PRESENTED) An apparatus to select a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet, the apparatus comprising:

a plurality of queues configured to store packets currently awaiting processing, each queue associated with one of the coprocessors, each queue associated with a cumulative load that represents a load to process packets in that queue;

a data structure configured to store processing rates, each processing rate associated with one of the coprocessors; and

a processor configured to determine a size of the received packet, and in response to the size of the received packet, and the processing rate of each coprocessor, determine a cost to perform a processing operation on the received packet at each coprocessor, the processor further configured to combine the cost at each coprocessor with the cumulative load at that coprocessor to create an anticipated load at each coprocessor, and to select a particular coprocessor to perform the processing operation on the received packet in response to comparison of the anticipated load at each coprocessor.

32. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the processor is further configured to select a coprocessor with minimum anticipated load to perform the processing operation on the received packet.

33. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the processor is further configured to determine if congestion is present at an output port associated with the received packet, and if congestion is present, select a coprocessor with non-minimum anticipated load to perform the processing operation on the received packet.

34. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the cumulative load associated with each coprocessor is determined in response to sizes of packets awaiting processing in the queue associated with that coprocessor and the processing rate of that coprocessor.

35. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the processing operation is an encryption operation.

VV. EVIDENCE APPENDIX

None.

VVI. RELATED PROCEEDINGS APPENDIX

None.